each specific interim-fast-access-time node of the plurality of interim-fast-access-time nodes to the respective second range of the LBAs.

21. (original) A method according to claim 18, wherein the data is allocated into groups of data within the one or more slow-access-time-mass-storage nodes according to a pre-defined unit comprising an integral number of bytes of the data, and wherein the mapping comprises a correspondence between the interim-fast-access-time nodes and the groups of data.

22. (currently amended) A method according to claim 17 ~~for storing data, comprising:~~
~~storing the data in one or more slow-access-time-mass-storage nodes having respective~~
~~first ranges of logical block-addresses (LBAs);~~
~~assigning to each of a plurality of interim-fast-access-time nodes, configured to operate~~
~~independently of one another, a respective second range of the LBAs;~~
~~coupling the plurality of interim-fast-access-time nodes to receive data from and provide~~
~~data to the one or more slow-access-time-mass-storage nodes having LBAs within the respective~~
~~second range;~~
~~receiving input/output (IO) requests from host processors directed to specified LBAs; and~~
~~directing all the IO requests to the interim-fast-access-time node to which the specified~~
~~LBAs are assigned;~~
wherein the one or more slow-access-time-mass-storage nodes comprise one or more disks, and wherein the interim-fast-access-time nodes comprise random access memories.

23. (original) A method according to claim 17, wherein the plurality of interim-fast-access-time

8

nodes comprise respective location tables, wherein each location table comprises locations of the second range of the LBAs assigned to the respective interim-fast-access-time node.

24. (original) A method according to claim 17, wherein the respective second ranges are spread sufficiently evenly and finely so as to generate well-balanced loading for the plurality of interim-fast-access-time nodes.

25. (original) A method according to claim 17, wherein each of the plurality of interim-fast-access-time nodes are at an equal hierarchical level.

26. (original) A method according to claim 17, wherein the respective second ranges of the LBAs do not overlap.

27. (currently amended) A method ~~according to claim 17~~ for storing data, comprising:

storing the data in one or more slow-access-time-mass-storage nodes having respective first ranges of logical block addresses (LBAs);

assigning to each of a plurality of interim-fast-access-time nodes, configured to operate independently of one another, a respective second range of the LBAs;

coupling the plurality of interim-fast-access-time nodes to receive data from and provide data to the one or more slow-access-time-mass-storage nodes having LBAs within the respective second range;

receiving input/output (IO) requests from host processors directed to specified LBAs; and

9

Appln. No.: 10/620,249

directing all the IO requests to the interim-fast-access-time node to which the specified

LBAs are assigned;

wherein the plurality of interim-fast-access-time nodes comprise a first and a second

interim-fast-access-time node, and wherein at least some of the respective second ranges of the

LBAs of the first and the second interim-fast-access-time nodes comprise overlapping LBAs, so

that one of the first and second interim-fast-access-time nodes is operative as a redundant

interim-fast-access-time node.

28. (original) A method according to claim 17, wherein the one or more slow-access-time-mass-

storage nodes comprise a multiplicity of slow-access-time-mass-storage nodes and the respective

first ranges are spread sufficiently evenly and finely so as to generate well-balanced loading for

the multiplicity.

29. (original) A method according to claim 17, wherein the plurality of interim-fast-access-time

nodes comprises a first interim-fast-access-time node and a second interim-fast-access-time

node, and wherein the first and second interim-fast-access-time nodes have substantially equal

capacities.

30. (original) A method according to claim 17, wherein the plurality of interim-fast-access-time

nodes comprises a first interim-fast-access-time node and a second interim-fast-access-time

node, and wherein the first and second interim-fast-access-time nodes have different capacities.

31. (original) A method according to claim 17, wherein the plurality of interim-fast-access-time

10

nodes comprises a first interim-fast-access-time node and a second interim-fast-access-time node, and wherein the one or more slow-access-time-mass-storage nodes comprise a first slow-access-time-mass-storage node which is coupled to only receive data from and provide data to the first interim-fast-access-time node and a second slow-access-time-mass-storage node which is coupled to only receive data from and provide data to the second interim-fast-access-time node.

32. (original) A method according to claim 17, wherein the plurality of interim-fast-access-time nodes comprises a first interim-fast-access-time node and a second interim-fast-access-time node, and wherein the one or more slow-access-time-mass-storage nodes comprise a first slow-access-time-mass-storage node and a second slow-access-time-mass-sto- rage node which are coupled to receive data from and provide data to the first and the second interim-fast-access-time nodes.

33. (previously presented) A system for transferring data to and from one or more slow-access-time-mass-storage nodes which store data at respective first ranges of logical block addresses (LBAs), comprising:

a plurality of interim-fast-access-time nodes, configured to operate independently of one another, each interim-fast-access-time node being assigned a respective second range of the LBAs and coupled to receive data from and provide data to the one or more slow-access-time-mass-storage nodes within the respective second range; and

11

one or more interface nodes, which are adapted to receive input/output (IO) requests from host processors directed to specified LBAs and to direct all the IO requests to the interim-fast-access-time node to which the specified LBAs are assigned;

wherein the interim-fast-access-time nodes are configured to be reassignable to a further second range of the LBAs.

34. (original) A system according to claim 33, wherein the one or more interface nodes comprise a mapping between the interim-fast-access-time nodes and the LBAs, and wherein the one or more interface nodes are adapted to convert the IO requests to one or more requests and to direct the one or more requests to respective one or more interim-fast-access-time nodes in response to the mapping.

35. (original) A system according to claim 34, wherein the mapping comprises a function relating each specific interim-fast-access-time node of the plurality of interim-fast-access-time nodes to the respective second range of the LBAs.

36. (original) A system according to claim 34, wherein the mapping comprises a table relating each specific interim-fast-access-time node of the plurality of interim-fast-access-time nodes to the respective second range of the LBAs.

37. (original) A system according to claim 34, wherein the data is allocated into groups of data within the one or more slow-access-time-mass-storage nodes according to a pre-defined unit of the storage system comprising an integral number of bytes of the data, and wherein the mapping

12

Appln. No.: 10/620,249

comprises a correspondence between the interim-fast-access-time nodes and the groups of data.

38. (original) A system according to claim 33, wherein the one or more slow access time-mass-storage nodes comprise one or more disks, and wherein the interim-fast-access-time nodes comprise random access memories.

39. (original) A system according to claim 33, wherein the plurality of interim-fast-access-time nodes comprise respective location tables, wherein each location table comprises locations of the second range of the LBAs assigned to the respective interim-fast-access-time node.

40. (original) A system according to claim 33, wherein the respective second ranges are spread sufficiently evenly and finely so as to generate well-balanced loading for the plurality of interim-fast-access-time nodes.

41. (original) A system according to claim 33, wherein each of the plurality of interim-fast-access-time nodes are at an equal hierarchical level.

42. (original) A system according to claim 33, wherein the respective second ranges of the LBAs do not overlap.

43. (previously presented) A system for transferring data to and from one or more slow-access-time-mass-storage nodes which store data at respective first ranges of logical block addresses (LBAs), comprising:

13

Appln. No.: 10/620,249

a plurality of interim-fast-access-time nodes, configured to operate independently of one another, each interim-fast-access-time node being assigned a respective second range of the LBAs and coupled to receive data from and provide data to the one or more slow-access-time-mass-storage nodes within the respective second range; and

one or more interface nodes, which are adapted to receive input/output (IO) requests from host processors directed to specified LBAs and to direct all the IO requests to the interim-fast-access-time node to which the specified LBAs are assigned;

wherein the plurality of interim-fast-access-time nodes comprise a first and a second interim-fast-access-time node, and wherein at least some of the respective second ranges of the LBAs of the first and the second interim-fast-access-time nodes comprise overlapping LBAs, so that one of the first and the second interim-fast-access-time nodes is operative as a redundant interim-fast-access-time node.

44. (original) A system according to claim 33, wherein the plurality of interim-fast-access-time nodes comprises a first interim-fast-access-time node and a second interim-fast-access-time node, and wherein the first and second interim-fast-access-time nodes have substantially equal capacities.

45. (original) A system according to claim 33, wherein the plurality of interim-fast-access-time nodes comprises a first interim-fast-access-time node and a second interim-fast-access-time node, and wherein the first and second interim-fast-access-time nodes have different capacities.

46. (original) A system according to claim 33, wherein the plurality of interim-fast-access-time

14

nodes comprises a first interim-fast-access-time node and a second interim-fast-access-time node, and wherein the one or more slow-access-time-mass-storage nodes comprise a first slow-access-time-mass-storage node which is coupled to only receive data from and provide data to the first interim-fast-access-time node and a second slow-access-time-mass-storage node which is coupled to only receive data from and provide data to the second interim-fast-access-time node.

47. (original) A system according to claim 33, wherein the plurality of interim-fast-access-time nodes comprises a first interim-fast-access-time node and a second interim-fast-access-time node, and wherein the one or more slow-access-time-mass-storage nodes comprise a first slow-access-time-mass-storage node and a second slow-access-time-mass-sto- rage node which are coupled to receive data from and provide data to the first and the second interim-fast-access-time nodes.

48. (previously presented) A method for transferring data to and from one or more slow-access-time-mass-storage nodes which store data at respective first ranges of logical block addresses (LBAs), comprising:

assigning to a plurality of interim-fast-access-time nodes, configured to operate independently of one another, respective second ranges of the LBAs;

coupling the plurality of interim-fast-access-time nodes to receive data from and provide data to the one or more slow-access-time-mass-storage nodes having LBAs within the respective second ranges;

receiving input/output (IO) requests from host processors directed to specified LBAs; and

15

directing all the IO requests to the interim-fast-access-time node to which the specified

LBAs are assigned;

wherein the interim-fast-access-time nodes are configured to be reassignable to a further

second range of the LBAs.

49. (previously presented) A method according to claim 48, wherein the IO requests are

directed to one or more interface nodes, wherein the one or more interface nodes comprise a

mapping between the interim-fast-access-time nodes and the LBAs, and wherein the one or more

interface nodes are adapted to convert the IO requests to one or more LBA requests and to direct

the one or more LBA requests to respective one or more interim-fast-access-time nodes in

response to the mapping.

50. (original) A method according to claim 49, wherein the mapping comprises a function

relating each specific interim-fast-access-time node of the plurality of interim-fast-access-time

nodes to the respective second range of the LBAs.

51. (original) A method according to claim 49, wherein the mapping comprises a table relating

each specific interim-fast-access-time node of the plurality of interim-fast-access-time nodes to

the respective second range of the LBAs.

52. (original) A method according to claim 49, wherein the data is allocated into groups of data

within the one or more slow-access-time-mass-storage nodes according to a pre-defined unit

comprising an integral number of bytes of the data, and wherein the mapping comprises a

16

correspondence between the interim-fast-access-time nodes and the groups of data.

53. (original) A method according to claim 48, wherein the one or more slow-access-time-mass-storage nodes comprise one or more disks, and wherein the interim-fast-access-time nodes comprise random access memories.

54. (original) A method according to claim 48, wherein the plurality of interim-fast-access-time nodes comprise respective location tables, wherein each location table comprises locations of the second range of the LBAs assigned to the respective interim-fast-access-time node.

55. (original) A method according to claim 48, wherein the respective second ranges are spread sufficiently evenly and finely so as to generate well-balanced loading for the plurality of interim-fast-access-time nodes.

56. (original) A method according to claim 48, wherein each of the plurality of interim-fast-access-time nodes are at an equal hierarchical level.

57. (original) A method according to claim 48, wherein the respective second ranges of the LBAs do not overlap

58. (previously presented) A method for transferring data to and from one or more slow-access-time-mass-storage nodes which store data at respective first ranges of logical block addresses (LBAs), comprising:

17

assigning to a plurality of interim-fast-access-time nodes, configured to operate independently of one another, respective second ranges of the LBAs;

coupling the plurality of interim-fast-access-time nodes to receive data from and provide data to the one or more slow-access-time-mass-storage nodes having LBAs within the respective second ranges;

receiving input/output (IO) requests from host processors directed to specified LBAs; and

directing all the IO requests to the interim-fast-access-time node to which the specified LBAs are assigned;

wherein the plurality of interim-fast-access-time nodes comprise a first and a second interim-fast-access-time node, and wherein at least some of the respective second ranges of the LBAs of the first and the second interim-fast-access-time nodes comprise overlapping LBAs, so that one of the first and second interim-fast-access-time nodes is operative as a redundant interim-fast-access-time node.

59. (original) A method according to claim 48, wherein the plurality of interim-fast-access-time nodes comprises a first interim-fast-access-time node and a second interim-fast-access-time node, and wherein the first and second interim-fast-access-time nodes have substantially equal capacities.

60. (original) A method according to claim 48, wherein the plurality of interim-fast-access-time nodes comprises a first interim-fast-access-time node and a second interim-fast-access-time node, and wherein the first and second interim-fast-access-time nodes have different capacities.

18

Appln. No.: 10/620,249

61. (original) A method according to claim 48, wherein the plurality of interim-fast-access-time nodes comprises a first interim-fast-access-time node and a second interim-fast-access-time node, and wherein the one or more slow-access-time-mass-storage nodes comprise a first slow-access-time-mass-storage node which is coupled to only receive data from and provide data to the first interim-fast-access-time node and a second slow-access-time-mass-storage node which is coupled to only receive data from and provide data to the second interim-fast-access-time node

62. (original) A method according to claim 48, wherein the plurality of interim-fast-access-time nodes comprises a first interim-fast-access-time node and a second interim-fast-access-time node, and wherein the one or more slow-access-time-mass-storage nodes comprise a first slow-access-time-mass-storage node and a second slow-access-time-mass-sto- rage node which are coupled to receive data from and provide data to the first and the second interim-fast-access-time nodes.

63. (previously presented) A storage system according to claim 1, wherein the interim-fast-access-time nodes are configured to be reassigned by a management node.

64. (currently amended) A storage system according to claim 1, wherein the interim-fast-access-time nodes are configured to be reassigned based on a <u>removal</u> ~~failure~~ of <u>at least</u> one of ~~one or more~~ <u>the plurality of</u> interim-fast-access-time nodes <u>or</u> ~~and~~ one ~~or more~~ <u>of the plurality of</u> slow-access-time-mass-storage nodes.

19

Appln. No.: 10/620,249

65. (currently amended) A storage system according to claim 1, wherein the interim-fast-access-time nodes are configured to be reassigned based on rebalancing a load between ~~at least~~ one of ~~one or more~~ the plurality of interim-fast-access-time nodes or ~~and~~ one ~~or more~~ of the plurality of slow-access-time-mass-storage nodes.

20